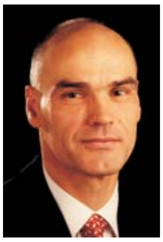

Directory Services: Definition – Status – Trends

Dr. Horst Walther

SiG Software Integration GmbH

Introduction

The market for Directory Services (DS) has expanded considerably over the past three years.

- *The previously so clear vision has become foggy:* Directory Services are fusing with Meta and Virtual DS, integration tools without DS emerging. A borderline towards EAI is required
- *Standards:* the LDUP initiative is shipwrecked. It is unclear how, and if, this hole in server-to-server communication will be filled. Standards and protocols from the Web Services field dominate the development
- *Directory enabled applications:* Resource Provisioning applications have entered the marketplace as a logical extension of identity and resource information, in order to process the necessary workflows. Complete user environments, like e.g. the mySAP infrastructure, become DS clients and, in connection with DS, offer comprehensive Identity Management functions
- *Vendors reposition themselves;* new vendors enter the market and some have merged

A different, immaterial (i.e. not material), development seems to be of particular consequence, however: a shift in perspective

- away from DS as isolated technical structural components, acquiring steadily increasing functionality but still seeking full appreciation as independent principal components of what is normally understood to be an enterprise structure
- towards regarding the process of Identity Management as a necessary administrative task in enterprises, without which e-business remains unsafe. Among the supporting components DS is probably the most important core component

This new orientation towards processes and thus the usefulness for individuals, and enterprises or other organisations, was overdue.

Three important inhibitors have so far been in the way of widespread acceptance and implementation of DS:

- The schism between ITU and IETF
- A lack of understanding of the usefulness of DS; and
- The difficulty of justifying infrastructure investments

The Schism Between ITU and IETF

The comprehensive *X.500* family of protocols was developed in the 1980s by CCITT and later by ITU. The results were comprehensive, well thought out specifications, only, however, implemented by a few organisations on the technology available at the time.

On top of this the standardisation process under the mainly governmental institutes appeared heavy, bureaucratic and suffocatingly slow to interested users.

In comparison, the IETF, in no way less successful, made due without costly standardisation rounds, coordinated concepts across the Internet through the RFC process and implemented only what was necessary at the time. This resulted in the LDAP protocol, version 3 of which is currently widely implemented. Later, LDAP comprised a family of protocols and extensions, in fact a full-blown Directory Service.

However, commercial success crippled the life of development. The full X.500 functionality was neither finished as specified nor re-invented and developed independently. Holes in functionality prevented both the physical implementation of different DS next to each other (more or less due to the lack of the equivalent of DSP / DISP in the LDAP world), and the understanding of the nature of DS.

Lacking Understanding of DS Usefulness

DS have really never been fully understood. This is of course a provocative statement. It is nonetheless correct to say that an understanding of the functionality and sensible limitations of the application of DS was the preserve of a small group of professionals.

E.g., it was difficult to make it clear to the majority of computer scientists why you couldn't simply use one of the proven relational RDBMS'es instead of a DS, irrespective of whether it was built around X.500 or the LDAP family.

Also at the technical level the positioning of DS as a further type of specialisation of data base systems was missing. Following the development of OLTP database systems for transaction processing and OLAP database systems for the analysis of pre-concentrated high-volume data, a DS is a special DBMS, optimised for frequent read processes involving (short) individual records[1].

Justifying Infrastructure Investments

In some cases DS are supplied free of charge. Other vendors charge per seat, leading to high costs in large organisations. In both cases the enterprise-wide introduction affects many departments. The project cost is, because of the multitude of processes needing to be co-ordinated, normally high; or *very* high.

Because a DS, due to its property as an infrastructure component, only creates value indirectly, it was initially very difficult to justify investments in DS.

Only the realisation that DS as storage for identities form the necessary foundation for Identity Management, and that it is neither safe nor efficient to create an e-business without IM, delivers the arguments that top management requires.

There are a number of reasons for this realignment of thought:

- *Increasing turbulence:* change has become the norm. Workers perform a particular role in the business for shorter periods than earlier. They switch between departments, work in project groups or move to a subsidiary for a few weeks
- *Increasing IT penetration:* office work nowadays almost exclusively involves the use of IT resources like PCs, e-mail and company intranets
- *Increased security awareness:* experience of the dangers of the Internet, the high level of dependency on IT and not least current global events have lead to a higher security awareness. "Can I borrow your password for a mo'?" is no longer accepted
- *External connections:* the electronic chaining of business processes into an on-line business introduces risks that require suitable governance. E.g., banks have to conform to the Basel II rules requiring that they set aside reserves to cover the operational risks associated with their internal business processes. These can subsequently only be lowered, according to an internal rating model, when it can be demonstrated to the regulatory authority that current business processes engender lower risks than corresponding to the associated reserves

Professionalising Identity Management

Just as for DS, the concept *Identity Management* (IM) is still neither widely understood nor precisely defined. For this reason we have chosen Microsoft's definition (see *DS' Role in Identity Management* next page), which to us best describes the concept of Identity Management best.

In 2002 the Open Group had made an attempt, *Business Scenario: Identity Management*, to connect the use of DS and other IM concepts with business scenarios [1]. The focus in this "Issue 1" from the Open Group is, however, quite narrowly conceived, the depth of reasoning in the documentation low.

The impulse towards a professionalised approach to IM was caused by the emergence of the controversial Microsoft product *Passport*. Passport was, as implied by the name, intended

[1] Example, CNN interactive: The implementation of an RDBMS did not deliver the required performance. By using a DS (eDirectory by Novell) CBB was able to guarantee access times below 250 ms at more than 2000 requests per second, corresponding to more than 5 millions per day.

November 2001, *Should I Use a Directory, a Database, or Both?*, Vikas Mahajan to prepare for a universal electronic identity. This, in itself welcome, initiative was met by strong criticism by a part of the market, because this, like issuing passports or personal IDs, was seen as the exclusive domain of the state, not to be usurped by a highly commercial company.

The *Liberty Alliance* was founded as a result of this, spearheaded by the competitor SUN Microsystems. They did not deliver a product similar to Passport but a set of specifications for the implementation, agreed upon by their members. A strategic goal clearly aimed at Passport has emerged from the current specifications.

Currently two more organisations work to take this discipline further: OASIS could be seen as being in close alignment with the Liberty Alliance, providing the more advanced technical underpinnings while the WS-Federation initiative mainly driven by IBM and Microsoft is producing a competing all-encompassing approach but is still well behind Liberty.

It is expected that these efforts will eventually lead to an Identity Management system spanning across enterprises, *Federated Identity Management*.

Classification of Directory Services

The Core Functionality of Directory Services

A definition which is currently universally accepted goes like this (see Figure 1):

A *Directory Service* is a specialised database system, usually used to store information about persons and associated resources. A DS is optimised to allow users and applications read access to retrieve short individual records. It is thus clearly distinguished from transaction database systems (OLTP) associated with a high frequency of changes and a low number of records; and from data warehouse systems (OLAP) with a low frequency of changes and a considerable amount of results to store.

In the current document it is assumed that DS in a more narrow sense conform with X.500 or at least allow access per LDAP version 3.

A *Directory* is the collection of information managed by means of this DS.

DS' Role in Identity Management

Today DS play an important role as infrastructure components for Identity Management. IM processes are generally well understood. In this context we understand IM as the *complete processing of digital personal identities* [2].

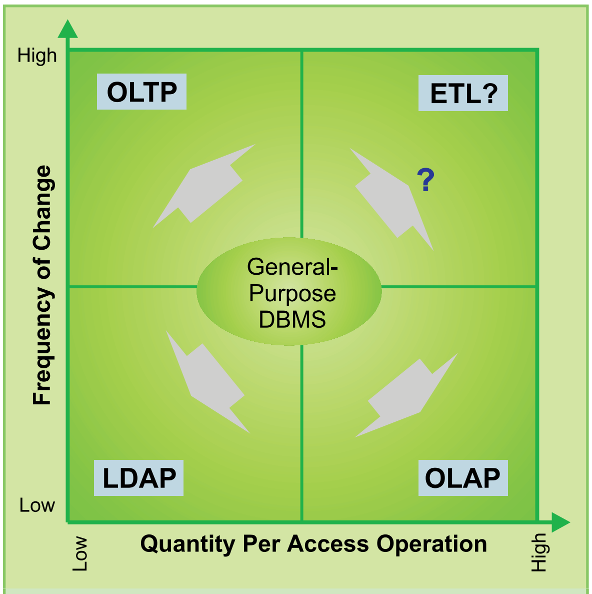

Figure 1 - **Database Categories**

The information, as well as the business processes needing it, are, in current organisations, normally both organisationally and from the point of view of technical support, distributed and uncoordinated in terms of implementation. The idea of a unified presentation of this information can thus be interpreted as a new management layer.

The information units in a digital identity comprise the attributes which identify a person:

- Business function
- Role
- Location
- Access permissions to business resources; and
- Personal information such as salary, bonuses and length of employment

IM can be arranged into three groups that serve the processing of digital identities (Figure 2):

- *Identity Administration:* management of digital identities, their relation to organisational departments and their access to resources
- *Community Management:* authenticating, preparing / publishing and authorising persons according to their digital personal identities
- *Identity Integration:* mechanisms to attain synchronisation and actualisation of digital identi-

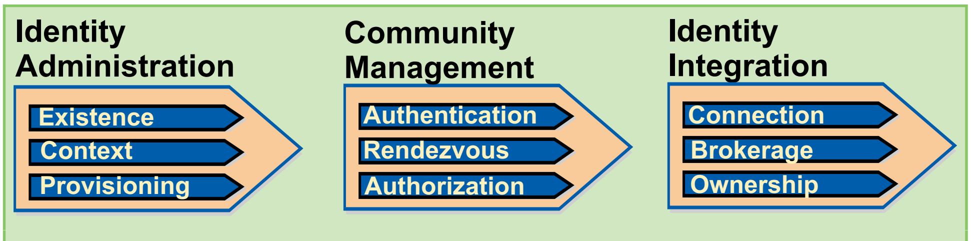

Figure 2 - **Indentity Management**

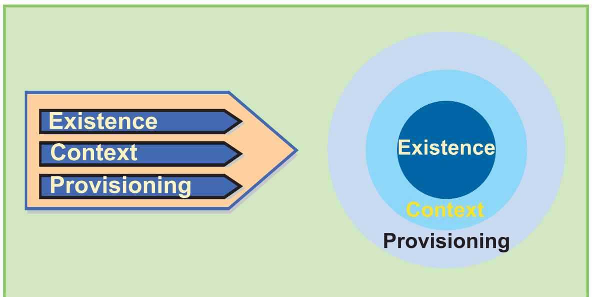

Figure 3 - **Identity Administration**

ties that are distributed across the organisation and contain partially redundant information

Identity Administration
The process of Identity Administration can be organised into three spherical groups, each on top of the previous, like an onion (Figure 3).

- *Existence:* creation, management and synchronisation of digital identities
- *Context:* management of the relationship of the person to the organisation (role) and his resources (permissions)
- *Provisioning:* supplying the person with the resources corresponding to his role and establishing such permissions in the target systems as required to control his resource usage

The *core* is formed by the processes required to manage unique digital identities.

The *middle sphere* contains the processes that manage the company roles; the connection between roles and permissions to resource access; and the extension of roles to digital identities.

The *outer sphere* contains the processes that supply persons with resources, either defined by his role, or additional resources.

The purpose of DS in the context of IM is supporting the process groups *Existence* and *Context*.

Development of Directory Services
Directory Services in the shape described here hark back to CCITT's (later ITU's) X.500 standards. Later developments, however, have largely been guided by technological influences.

In the initial phase of commercial Internet exploitation the fulfilment of X.500 requirements (OSI protocol stack, .1, etc) was too demanding for the hardware available. As a result a lightweight version of the X.500 access protocol, DAP, later called LDAP, developed.

Subsequently hardware performance became less and less of a bottleneck. In particular, the availability of cost-effective and powerful database servers, in conjunction with the fact that most of the identity information was already stored in non-LDAP repositories, led to the development of so-called *Virtual Directory Services*. These purposely renounced read access optimising. The data access by a requester was only simulated, while in reality the original data source was accessed.

A further development, the increasing performance of both public and private networks, led to a reduction in the demand for the X.500 replication protocols, DSP and DISP, necessary to operate distributed databases. Instead it became possible to access central databases directly.

This development could have led to unified and centrally operated enterprise data storage repositories. However, so far the stage of unified enterprise-wide storage has not been realised by established corporations. There are reasons to believe that this situation will not change soon.

E.g., this has led the Gartner Group to predict that until 2006, 80% of mid-sized and large enterprises will fail to standardise on a single and unified enterprise-wide database [3].

Database Consolidation
DS already possess in-built integration capabilities. Three material properties are a measure of this:

- It is possible to navigate the hierarchical structure tree. I.e., it is not necessary for objects to be known in order for them to be located
- Parts of the schema are standardised. Certain schema extensions are used in addition to the Attribute Types (X520) and Object Classes (X521) defined in the X.500 standard, e.g. inetOrgPerson (RFC2798)
- A well defined and universally used access protocol is available via LDAP

Nevertheless, powerful conservative and dividing forces oppose this normative effect, which in large measure can be applied to staff databases in the individual corporations.

The powers, which work against the target of consolidating databases, are too strong to be conquered in most companies. They are:

- *Platform dependencies:* Certain operating systems, e.g. NetWare and Windows 200x, are delivered with a DS. This DS is tied to the operating system and cannot be replaced by a different one
- *Explicit application dependencies:* Certain business applications are also tightly connected to a particular DS. As in the case with platform dependency it is not possible to substitute a DS of choice
- *Implicit application dependencies:* Many application vendors explicitly support one or more DS and make claims such as "supports any LDAP v.3 SD", without being able to furnish any guarantees

What is LDAP?

LDAP stands for *Lightweight Directory Access Protocol*. LDAP is a client-server protocol for accessing a directory service. It was initially used as a front-end to X.500 directory services, but can also be used with stand-alone and other kinds of directory servers.

LDAP was developed by a group of developers led by Tim Howes at the University of Michigan in the early 1990s as a low-cost way for PCs to access complex directories that were based on the X.500 global directory standards. X.500 was developed by the *Comité Consultatif International Téléphonique et Télégraphique* (CCITT) and later the *International Telephone Union* (ITU); it was published as a set of international standards by the *International Organization for Standardization* (ISO).

X.500, published in 1988 and updated in 1993, focuses on communication between a client and a server, specifying how information is stored and accessed in a large directory or in multiple interconnected directories forming global directory services.

While the X.500 standards were comprehensive and well engineered with respect to a homogeneous and standards-compliant technology, implementation proved to be cumbersome and overhead intensive; this delayed widespread adoption.

LDAP emerged to bridge the gap between the end user and global directories at a cost of just about 15% loss of functionality compared to the corresponding X.500 protocol, DAP (Directory Access Protocol); hence a "lightweight" DAP. The Development was driven by the *Internet Engineering Task Force* (IETF) and backed by an urgent demand during the early phase of the commercialisation of the Internet. LDAP quickly became the solution of choice for all types of directory services applications running over the Internet Protocol (IP).

LDAP has undergone a rapid evolution within the IETF. It is documented in a set of Requests for Comments (RFCs), e.g.:

- RFC 1487 (July 1993) LDAP v1
- RFC 1777 (March 1995) LDAP v2
- RFC 2251 (December 1997) LDAP v3

Version 3 of LDAP has improved on earlier versions by increasing security (using *Simple Authentication and Security Layer* – SASL and encryption via the *Secure Socket Layer* – SSL). Version 3 also enables one LDAP server to make a referral to one or more other LDAP server(s); it allows for additional services (extensions); enables integration with multiple vendor databases; and it allows the use of non-ASCII and non-English characters to gain a true international flavour.

The evolution of LDAP has slowed down considerably today. Nevertheless some useful extension have reached the status of Internet-Drafts under the umbrella of the LDAP (v3) Revision (ldapbis) working group:

- LDAP: String Representation of Distinguished Names
- LDAP: The Protocol
- LDAP: String Representation of Search Filters
- LDAP: Authentication Methods and Connection Level Security Mechanism
- LDAP: Uniform Resource Locator
- LDAP: Syntaxes and Matching Rules
- LDAP: Technical Specification Road Map
- LDAP: Directory Information Models
- LDAP: Internationalized String Preparation
- IANA Considerations for LDAP

LDAP enables one to "locate organizations, individuals, and other resources such as files and devices in a network, whether on the Internet or on a corporate intranet," whether or not the domain name, IP address, or geographic whereabouts are known.

An LDAP directory entry is a collection of attributes with a name, called a *Distinguished Name* (DN). The DN refers to the entry unambiguously. Each of the entry's attributes has a type and one or more values. The types are typically mnemonic strings, like "cn" for common name, or "mail" for e-mail address. The values depend on the type. For example, a mail attribute might contain the value "donald.duck@disney.com". A jpeg-Photo attribute would contain a photograph in binary JPEG/JFIF format.

LDAP directory entries are arranged in a hierarchical structure that reflects political, geographic, and/or organisational boundaries. Entries representing countries appear at the top of the tree. Below them are entries representing states or national organisations. Below them might be entries representing people, organisational units, printers, documents, or just about anything else.

Many current email clients, including Microsoft Outlook, Eudora, and Netscape Communicator, use some form of LDAP database to look up email addresses. Internic and Infospace are two examples of big public look-up services built with LDAP.

Some of these solution providers, Sun and Microsoft in particular, have designed JNDI and ADSI APIs so that you can connect with any kind of directory service. This is akin to what ODBC or JDBC is to an RDBMS.

- *Suitability for specific tasks:* DS products have different strong and weak points. They use different replication mechanisms, have different levels of error tolerance, adhere to different security philosophies and have led to different expectations regarding scalability

Thus, enterprises are often forced to run several DS in parallel. For this reason the use of DS in the growing user landscape of large and medium sized enterprises carries a high integration burden. A means of delivering the integration advantages lies in the establishment of Meta Directory Services and/or Virtuel Directory Services.

Meta Directory Services

What is a Meta Directory Service (MDS)?

The concept of an MDS is not precisely defined and often associated with different functionality.

In a sense the name is an unfortunate choice because it creates the association that this is where meta-data is stored. The Greek word *meta* means *above*, hence the (unexplainable) system of rules governing what is above the physical world came to be know as metaphysics. In the same way metadata contains data about other data, i.e. a *description* of the data.

The distinctive functionality of an MDS, however, lies in the ability to integrate; in particular by combining information from different databases. Thus, MDS should really be called *Database Integration Services*, i.e. services that integrate data from several data repositories, a more appropriate name.

Directory-agnostic approach: when it comes to the practical implementation in companies it is often the case that source, and sometimes even target, databases fail to be prepared by means of read-optimised LDAP Directory Services. In these cases the MDS is reduced to its essential core, its integration and data processing functionality. This functionality in its purest form is represented by the IBM Directory Integrator (by the Norwegian company Metamerge, acquired by IBM).

The inclusion of non-LDAP sources into the synchronisation of identity information can only be seen as a sign of the maturity and acceptance of Identity Management efforts. In reality, most authoritative sources of identity information are not kept in LDAP or X.500 DS, so their integration is more important than the synchronisation of installed LDAP or X.500 DS. This synchronisation can be brought about, as a side effect, by a high-performance meta-component. In this case the implementation of a native, or a standard inter-server communication protocol such as DSP or DISP from the X.500 group, or LDUP – the separate attempt by the IETF, is not required. As a consequence the perhaps most important reason for implementing X.500 products in mature enterprise systems falls away.

*MDS Componen*ts

The most important components in an MDS are defined as:

- *Join-Processing:* a rule-based conversion between different representation formats of fields with identical semantics in different directories
- *Event triggering:* reacting to changes in a data constellation or other events by launching a simple work flow
- *Connectors:* modules allowing access to non-standard data sources (data sources that cannot be accessed by means of DAP, LDAP or ANSI-SQL)

Conspicuously, the DS used (optionally) to store the resulting information is missing from this list. This is a, in some cases (Virtual DS) dispensable, component, which contributes little to the integration work and which can often easily be added by a third party supplier.

The Functionality of MDS

Products marked under the label Meta Directory Services differ considerably in the manner in which they deliver this functionality.

Most MDS products offer multi-directional synchronisation between directories containing user information.

Many MDS products offer a unified view of the information they synchronise.

Some MDS products do not actually synchronise but only offer a unified view of the underlying data. These products are often sold under the label Virtual Directory Services.

A few MDS products are really middleware infrastructure products, enabling further general integration services.

In this paper we consider products having the first property mentioned above (synchronisation) and optionally the second (unified view) genuine MDS products, and we define a MDS as software delivering multi-directional information exchange between directories and other collections of user information [4].

Reasons for Implementing MDS

Many companies face the challenge of increasing the pace with which new applications are brought on-line.

At the same time important information about customers or business partners is often spread across numerous and often incompatible IT systems.

Whereas these types of tasks traditionally have been solved through development of interface programs, this option is mostly no longer available for reasons of cost and time limitations.

So far, no single vendor or architecture is able to completely fulfill the requirements of an enterprise infrastructure. The dream of a single unified directory from a single supplier, capable of providing unified access mechanisms, security, business functionality and information throughout the enterprise, has proven elusive. Nor does it look as if this will emerge any time soon.

Current efforts at standardisation have so far not delivered the necessary capabilities with regard to integration of independent enterprise systems.

As it is, systems have not been integrated in any useful fashion. Consequences can be that mission critical systems fail to scale, or that the extended business environment including partners, customers and suppliers cannot efficiently be made available.

Critical business agility increasingly depend on increasing the agility of the business infrastructure. The integration of legacy systems has become an important factor for scalability, agility and reduction of implementation costs. This can be solved by means of MDS.

Virtual Directory Services (VDS)

What is a VDS?
A is a service used to extract information from other DS and from miscellaneous other data sources. The moniker 'virtual' implies that a is not itself a DS but assumes the functionality of one. Thus, it produces a virtual representation of distributed data during read access. The rule-based consolidation of these records, stored in different directories (the so-called *join*), is saved by the MDS inside its own storage space.

In other words, VDS here also offer a real-time join. From the point of view of the read process the VDS simulates access to a directory, whereas in reality – somewhat similar to reading a view in a RDBMS – the original data sources are accessed. Compared to the situation some years ago modern fast hardware, no longer the bottleneck, enables this transformation from the original source data in real-time (in connection with intelligent caching) without any need for the added complexity of intermediate data storage.

VDS is a subset of MDS. Their special properties are:

- VDS don't synchronise underlying data but offer a unified view of information, synchronising data in the view
- They transform LDAP queries into queries on the original data sources, receive replies and present these to the requester in a unified view
- Since VDS, as opposed to DS, access the original data sources, these are not necessarily stored in a manner optimised for reading. For this reason the timing is determined by the latency of the data sources, in fact, of the slowest of them
- VDS permit the ownership of data to stay with departments (no political trench-fighting)
- The real strength, that VDS doesn't require additional physical storage, may also be interpreted as a disadvantage: the weakest system determines the reliability of the total system. Intelligent caching can mitigate this defect but this is again a step back in the direction of traditional MDS

Reasons for Implementing VDS
Implementing VDS is indicated in cases where:

- existing data repositories (SQL databases, etc) must be LDAP enabled
- you want to use them as firewall proxies: using VDS to reach data through firewalls can deliver a subset of the enterprise information to external users
- corporate data consolidation must be achieved without risking political rejection
- you need a proxy for schema and name-space transformations
- when real-time access is required (e.g. to support a PKI)
- you need to secure e-mail: in this case can be used to deliver a reduced set of address information on an address list. E.g. the case when classified documents must be sent encrypted but where not all possible recipients in an address book possess a qualified certificate

Interoperability

An Overview of Activities
There are currently three groups that work at furthering the development of standards in the area of directory services:

1. *IETF*, The Internet Engineering Task Force

    - *LDAPv3 Revision* (LDAPbis): one workgroup; reworking of the LDAPv3 core specifications (RFCs 2251-2256 and 2829-2831); has already matured to a draft standard
    - *LDAP Duplication / Replication / Update Protocol* (LDUP): the task of this group is to standardise the Master-Slave and the Multimaster LDAPv3 replication; it has taken over the LDAP Extensions workgroup

2. *OASIS*, the Organization for the Advancement of Structured Information Standards

- *Directory Services Markup Language* (DSML); version 2.0 is available. DSML is mainly a transposition of LDAP to XML (with extensions to allow batch queries and delete operations)

- *Security Assertion Markup Language* (SAML); XML-based standard intended to allow authentication and authorisation of queries across security domain boundaries; released in version 1.1; used by the Liberty Alliance as their basic protocol

- *Extensible Access Control Markup Language* (XACML 1.0): an XML specification for policies allowing access to information across the Internet; the result may be standardised access controls

3. *The Open Group*

- *Directory Interoperability Forum* (DIF): a forum, which is active world wide, and has set as its goal to develop open and interoperable directory services, and to further the implementation of these in industry; the OPEN group also issues the LDAP 2000 compatibility certificate

Discussion
It looks as if the IETF, with the exception of a few enthusiastic individuals, has lost either the interest or the patience when it come to driving further development of LDAP.

The only remaining organisation that still retain sufficient motivation and resources is the Directory Interoperability Forum of the Open Group. They enabled the LDAP 2000 initiative and support the LDAP branding programme.

Op top of that it is impossible to deny that the development of LDAP is stagnating. No new standard has emerged for four years in the context of LDAP. In the same period the application of XML has gained strength as *lingua franca*, probably not a coincidence.

It has become completely obvious that XML has come to be seen as the means of data exchange for the future. At the same time attempts to tie in LDAP with XML have been abundant.

The definition of the *Directory Services Markup Language* (DSML), an XML dialect to map LDAPv3, is one such attempt. After this initiative (brought to live by the company Bowstreet) threatened to falter because of inherent limitations in DSML 1.0, version 2.0 is now ready. The reviews have become considerably more friendly, e.g. SUN Microsystems have transformed the specification into their XMLDAP Gateway.

Other XML standard developments, e.g. the *Security Assertion Markup Language* (SAML) or the *Extensible Access Control Markup Language* (XACML) may possibly extend LDAP and thus also DSML with communications capabilities. By doing so they partly fulfill the promises of the work on the LDAP Extensions (LDAPExt).

Observed from some distance these developments appear to be rearguard firefighting by the LDAP community. The functions of LDAP can also be taken over by XML and the *Simple Object Access Protocol* (SOAP).

Market observers assume that even the communication between different directory services will be run across an XML layer. That way direct standardised communication between the directory services would no longer be absolutely required.

The emphasis is moving from directory integration to application integration, as indicated by the efforts to integrate directory information via open XML-based standards undertaken by standardisation groups like the WS* group or the Liberty Alliance.

Another important factor resides in firewall technology as currently implemented. Enterprises usually allow communication into the Internet over port 80. XML data travel via HTTP through port 80. LDAP uses ports 389, 3268 / 636 and 3269, not normally allowed by firewall systems.

While this may not seem plausible from a security point of view it may actually be seen as an advantage that XML can be transferred through port 80 while LDAP often remains shut off behind the firewall.

Conclusions: Further development of LDAP is stagnating and will not lead to the results that were hoped for; technologies based on XML will replace or be inserted on top of LDAP; even if the importance of DS should increase, the development will no longer be in favour of LDAP; while LDAP will remain the working horse for quite a while the future will be directory services without LDAP.

Suppliers and Their Positioning

Changes Since 2000
As late as in 2000 the discussion was dominated by questions such as:

- when is the LDAP community going to deliver full X.500 functionality?

- will all future DS also be MDS?

Accordingly, all products we categorised according to these criteria:

- X.500 or LDAP

- DS or MDS

Subsequently some changes have emerged; the trends and developments that we find most important are listed below:

Bundling or unbundling: In 2001 the Gartner Group still prophesied that by 2005 DS would no longer be sold as stand-alone products but bundled with operating systems without additional charge. The reason given by the Gartner analysts was IBM's announcement that the product IBM SecureWay was to be given away for free [5]. As opposed to this Microsoft, with AD/AM, has taken a step in the opposite direction.

The re-emergence of Novell: Because of a steady development effort over the past 13 years, support for current technologies like XML and DSML, combination of LDAP and UDDI or cooperation with the Liberty Alliance, Novell has managed to win back previously lost confidence.

The hesitant acceptance of Microsoft's Active Directory (AD): The unexpectedly high stability of Windows 2000 made this a convincing system right from its introduction. For this reason market observers leaned towards the opinion that a major part of Windows 2000 users would also implement AD. Accordingly, Windows 2000 projects were started in many large corporations in 2000, which included as a central task the evaluation of implementation scenarios for AD. These revealed the complexity of such projects, leading to AD implementations being delayed or their scope being reduced. In the end only a small proportion of Windows 200x users have implemented Active Directory.

The emergence of new vendors (e.g. Radiant Logic, OctetString and MaXware) with implementation in modern technology and an emphasis on easier implementation-projects: these vendors address the main obstacle to the introduction of DS, the technical and, above all, organisational, complexity associated with its introduction. Additionally, they place particular emphasis on implementation aids and thus to a degree renounce individual implementation of DS. The result is Virtual Directory Services and meta-components (integration components) for DS and other sources of (and sinks for) data. They are typically developed in J2EE technology, use XML and message-queuing technologies and in no way position their products as substitutes for existing storage solutions but rather – with a degree of tactical understatement – as a kind of reporting tools and integration aids.

IBM persists: earlier, IBM was considered to be the taillight when it came to implementing and marketing directory services. IBM's strategy with regard to DS used to engender incredulity in market observers and players, or were frankly met with denial. However, according to all commentators IBM made a shrewd chess move by acquiring the Norwegian technology leader, the software house Metamerge, in a period of low demand and contrary to the business cycle, as well as the provisioning vendor Access360. By consequent integration of components and systems with relevance to Identity Management into the IBM Corporation IBM managed to move from an accepted position as the taillight, somehow past all the marketeers, into a tentative headlight position.

The drying out of IETF activities in the DS area: The LDUP initiative, supposed to produce a standardised inter-server replication protocol, was technically quite successful. It was, however, turned into a failure due to lack of vendor acceptance. Further development on LDAP has also ground to a halt. We shall probably never see an LDAPv4. The CEO of the Burton Group, Jamie Lewis, is even said to controversially have pronounced LDAP dead.

XML has become the standard-bearing technology: market experts have only reluctantly admitted that a potential LDAP, and of course also DAP, competitor has arrived with XML. Today, however, it is irrefutable that XML dialects dominate the development. LDAP is thus threatened to suffer the same destiny that earlier befell the X.500 access protocol DAP. After all, LDAP was only meant to be an access protocol for X.500 Directory Services that was easier to implement and consumed fewer resources. LDAPv3 introduced the first "LDAP Directory Services" emerging in the X.500 family without further borrowing from elsewhere. It is noteworthy that following the somewhat unfortunate beginnings with the LDAP-XML offspring DSML, directory access without further detours became possible using an XML dialect. The detour around an LDAP packaged inside XML, like that offered by DSML, was no longer required. The specification for UDDI, a DS for web services built entirely on XML and without any reference whatsoever to existing DS technologies included, looks very much as the writing on the wall.

Web services dominate the discussion: web services, with their business possibilities and associated technical challenges, are a current discussion theme. SOAP and XML are important elements of the web communications architecture. The DS for web services, UDDI, is based on these technologies. A point of criticism is that this is a re-invention of the wheel and that all the problems still inherent to UDDI have now been solved in the LDAP area. To some market observers it looks as if history is about to repeat itself: also LDAP was originally positioned against the X.500 protocol DAP. It took several years until LDAP got even close to the X.500 functionality. So it is not surprising that several vendors, prominently including Novell, shunned the implementation of UDDI servers as part of the DS technology they brought to market. Next to this they (out of necessity) bridged the gap between Identity Management and Web Services Management and thus technically between the LDAP world and the XML / SOAP world.

Microsoft Passport and the Liberty Alliance: Microsoft's push to deliver a portable identity definition and implementation with the Passport product, featuring the ability to realise single

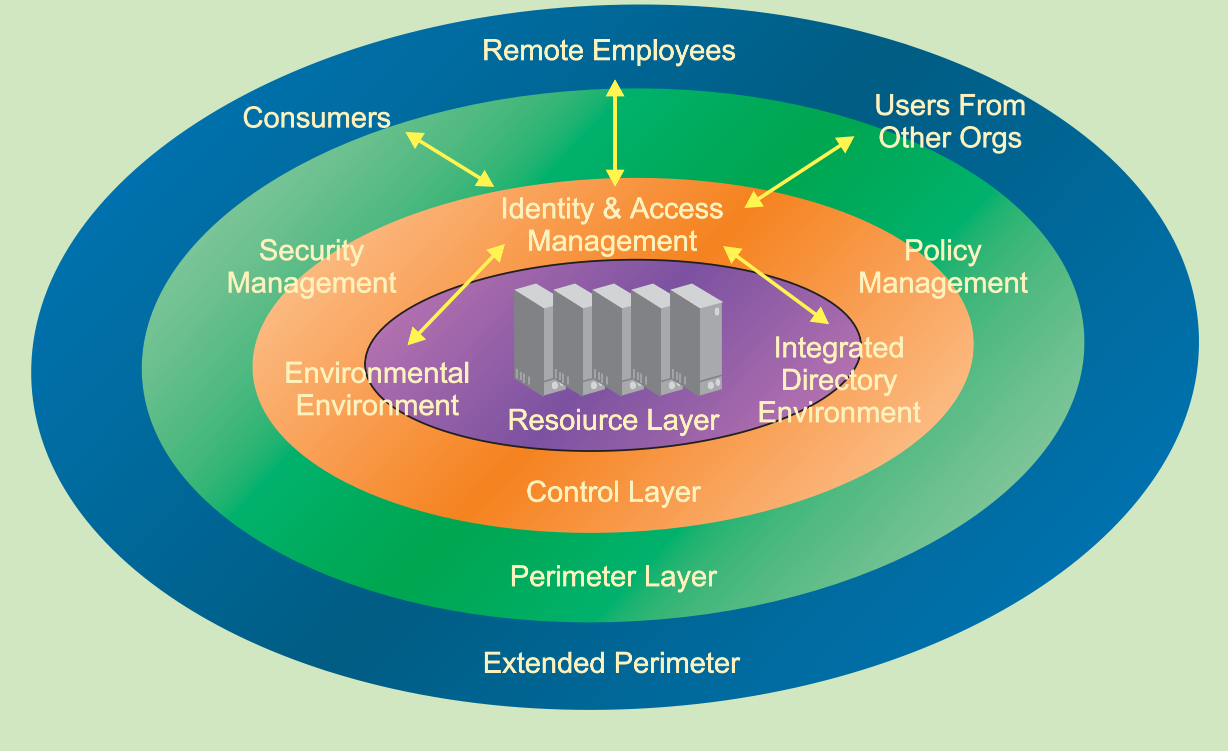

Figure 4 - **Multilayered VEN Security Architecture**

sign-on, particularly for B2C applications on the web (Web-SSO), led, after an initial bout of terror, to considerable discussion in the market and in the political arena. On the initiative of SUN Microsystems this resulted in the formation of the Liberty Alliance, which started to publish specifications for alternative products to Microsoft's offering. They decided that the product *Security Assertion Markup Language* (SAML) for web authentication should be the first. Integrating this into existing company infrastructures for Identity Management seemed obvious but this process is still not complete. The competition between the Liberty Alliance and the WS-Federation Group should not necessarily be seen as a disadvantage; to the contrary: it increases the probability of the near term delivery of acceptable results.

For the past half decade X.500 has been declared dead. Nonetheless, the X.500 products from Nexor, Siemens and Critical Path live on quite happily as stable niche products, without experts paying much attention, being critical components of the infrastructures of government and supra-national institutions, as well as some large enterprises. In 2000 Critical Path strengthened their position by acquiring their competitor and X.500 vendor PeerLogic and its product i500 Directory. In this niche we find a group of small companies, e.g. Data Connection Ltd, selling products of a quite astonishing functionality.

Trends

Convergence of the Network

The picture painted by Burton Group of the fusion of internal and external networks (Figure 4) may not emerge soon, being stopped by the force of existing structures and technologies. In the long term, however, it sounds realistic. The borders between Internet and intranet systems will disappear. The duality of trusted (internal) and untrusted (external) will be replaced by a new plurality of separate trust domains.

The fusion of networks requires a new security architecture. Technically the emphasis will be on interoperability, authentication and authorisation.

Identity Management Instead Of Directory Services

The long discussion about DS will fall silent. As in the past, it will not be possible to justify investment in infrastructure in the future without putting these in the context of a forceful business context proving necessity.

This business context seems to have emerged. It is called *Identity Management*. Purposes reaching across businesses and business categories such as the realisation of the promises of e-business; the increased dynamism in companies as mentioned above, with workers remaining in specific roles for shorter periods of time; and the general increase in security awareness deliver enough reasons for the 'unavoidability' of professional Identity management.

As a cross-organisational task IM thus acquires an importance similar to that of accounting or auditing in the corporate structure.

Add to that business-specific requirements. Regulatory requirements are imposed upon financial services suppliers in most countries, for example; consider also the new risk accounting measures included in Basel II, requiring operational risks to be taken into consideration when banks calculate their capital requirements.

Conclusions: The formerly rather under-employed DS will again spread – without anybody talking about it. They form the necessary foundation for Identity Management.

Identities Become Portable

The processes group *Identity Management*, the core of identity, has increased its value considerably as a result of some renewed attention, *in casu* the discussion about portable identity definitions (and products) raised by Microsoft Passport, and the corresponding initiatives in the Liberty Alliance and the WS-Federation group.

This contentious development will lead to the availability of a generally accepted portable electronic identity system (something no signature directive legislation has succeed in doing).

No credible alternative exists that can perform the technical conversion of this process group in IM using DS.

Conclusion: Portable identity definitions need DS as management systems. The specifications from the Liberty Group and/or the WS-Federation group will speed up this development very considerably.

Meta-Services Without Directory Services

While analysts two years ago had raised the provocative question, "Will all Directory Services be 'Meta' in the future?" – today Directory Services have instead reverted to their core capability, Identity Management.

Integration components, surrounding the DS like a ring, will become increasingly important. They will detach themselves from the DS, itself, to lead a life of their own as separate middleware tools.

The IBM Directory Integrator (ex Metamerge Integrator) personifies this type of next generation MDS and serves the competition as a template for functionality and operative properties.

Conclusion: Meta components will in the future complement the use of DS but as separate components, detached from the DS itself, without preference for a particular DS. They will become directory agnostic.

Convergence between Identity and (Web) Services Repositories

According to the Burton Group the next logical step is the convergence between LDAP DS and UDDI DS [6].

Companies like Novell, Sun and CA have implemented UDDI interfaces on their LDAP DS. BEA is moving in the opposite direction.

Conclusion: There may be enthusiastic opponents to the integration between identity and web services repositories. Nevertheless, much seems to indicate that LDAP and UDDI DS will merge.

References

[1] The Open Group, *Identity Management Business Scenario*, Issue 1, 15, July 2002

http://www.opengroup.org/downloads/bus-scenario-IM.pdf

[2] Microsoft, *Strategy White Paper: Enterprise Identity Management*

http://www.microsoft.com/windows2000/docs/_Toc456089786

[3] Enck J., *Enterprise Directories: Bet You Can't Deploy Just One*, June, 11, 2002

[4] Enck, J., *So Many Directories, So Little Consistency*, March 01 2002

[5] Enck, J., MacDonald, N., *Directory Services Market and Magic Quadrant*, 2001

[6] *UDDI And LDAP: A Perfect Fit?*, LDAPGuru, June 27, 2002

http://www.ldapguru.net/modules/news/article.php?storyid=157&PHPSESSID=4d9cd8bffb1033abc2ed8e33796ef07f

About the Author

Dr Horst Walther is CEO of SIG Software Integration GmbH in Hamburg, Germany.

Horst was awarded an M.Sc. in chemistry by Hamburg University after studying computer science, chemistry, sociology and oriental sciences. His thesis in the field of physical chemistry, centred around physical chemistry, plus additional work in technical chemistry and and computer science led to a Ph.D. from the same university.

Since 1975 Horst has been working in information technology, including being a consultant in different business areas, mainly banking and insurance, since 1984. His work includes IT system audits, development of IT strategies and as a specialty subject storage systems and their application in Identity Management. As a consultant to the German Savings Banks and individual institutions within this group he has conducted several research studies in the past few years, mainly regarding the roles of Directory Services in information management.